
Horizontal Integration: Broader Access Models for Realizing Information Dominance

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

20050126 016

MITRE

REPORT

Horizontal Integration: Broader Access Models for Realizing Information Dominance

JSR-04-132

Approved for Public Release

JASON
The MITRE Corporation
7515 Colshire Drive
McLean, Virginia 22102-7508
(703) 883-6997

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information estimated to average 1 hour per response, including the time for review instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE December 2004		3. REPORT TYPE AND DATES COVERED
4. TITLE AND SUBTITLE Horizontal Integration: Broader Access Models for Realizing Information Dominance			5. FUNDING NUMBERS 13049022-DC	
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The MITRE Corporation JASON Program Office 7515 Colshire Drive McLean, Virginia 22102			8. PERFORMING ORGANIZATION REPORT NUMBER JSR-04-132	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Defense Research and Engineering (ODDR&E) Director, Plans and Programs 3030 Defense Pentagon Room 3D108 Washington, DC 20301-3030			10. SPONSORING/MONITORING AGENCY REPORT NUMBER JSR-04-132	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release.			12b. DISTRIBUTION CODE Distribution Statement A	
13. ABSTRACT (Maximum 200 words) Horizontal integration refers to the desired end-state where intelligence of all kinds flows rapidly and seamlessly to the warfighter, and enables information dominance warfare.				
14. SUBJECT TERMS Information security			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	

Contents

1	INTRODUCTION	1
1.1	As Usual, Sun Tzu Had It Right	1
1.2	Today's Context Is Important for This Study	2
1.3	What Problems Need Solving?	4
2	THE PRESENT SYSTEM OF CLASSIFIED INFORMATION	7
2.1	The Present System Is Largely Unchanged Since the 1940s . .	7
2.1.1	Before 1940 There Was Only Informal Classification . .	7
2.1.2	The Present System Was Established in 1940 and Has Changed Little	9
2.2	In Principle, Risk Level Governs Classification	10
2.3	In Practice, Distribution Channels Are a Dominant Factor . .	12
2.4	Computers and Networks Present Enormous Challenges	13
2.5	There Is Increasing Evidence That the Present Construct Is Breaking Down	15
2.6	The Present System Has Some Good Constructs But Is Miss- ing Others	16
2.6.1	Constructs Present in the Present System	16
2.6.2	Missing Constructs in the Present System	18
3	TOWARDS A NEW SYSTEM FOR INFORMATION PROTECTION	21
3.1	A New System Should Satisfy Some Basic Criteria	21
3.2	The STU III Is a Good Example of Accepting Risk	22
3.3	IAD's "45 Day Study" Is Useful But Not Radical Enough . . .	23
3.4	Three Guiding Principles for Any New System	25
4	SPECIFIC PROPOSAL FOR A NEW PARADIGM	27
4.1	We Propose a Three-Phase Plan	27
4.2	Preparatory Phase (Phase 0)	27
4.2.1	Developing a Risk Model	28
4.2.2	Developing Necessary Infrastructure	31
4.2.3	"Enclaves" or Communities of Interest	33
4.2.4	NetTop Is an Important Technology	35
4.3	Risk is Tokenized in Phase 1	39

4.3.1	What is a Token?	39
4.3.2	How Are Tokens Denominated?	40
4.3.3	Phase 1 Tokens Are Not Fully Fungible	42
4.3.4	How Are Tokens Distributed?	43
4.3.5	How Are Information Producers Incentivized?	44
4.3.6	Steps Toward an Efficient Market Economy	46
4.4	Originator Control is Eliminated in Phase 2	48
4.4.1	Tokens Collapsed to a Few Broad Token Types	48
4.4.2	Within A Broad Token Type, Access Is Now Fungible	49
4.4.3	Phase 2 Requires Both a Risk Model and a Damage Model	50
4.4.4	Tokens Are Created And Distributed by a National Authority (Central Bank)	51
4.4.5	In a Tokenized System, Personnel Reliability Is a Con- tinuous Variable	52
5	SUMMARY AND CONCLUSIONS	57

1 INTRODUCTION

Horizontal integration refers to the desired end-state where intelligence of all kinds flows rapidly and seamlessly to the warfighter, and enables information dominance warfare.

1.1 As Usual, Sun Tzu Had It Right

Sun Tzu, while ignorant of modern remote sensing technology, well understood the basic principle of information dominance:

“O divine art of subtlety and secrecy! Through you we learn to be invisible, through you inaudible; and hence we can hold the enemy’s fate in our hands. ... By discovering the enemy’s dispositions and remaining invisible ourselves, we can keep our forces concentrated, while the enemy’s must be divided.”

Sun Tzu’s said this about horizontal integration:

“With none in the whole army should closer relations be maintained than with spies. None should be more liberally rewarded.”

But Sun Tzu is also direct regarding the importance of keeping secrets, both within the collector and user communities:

“If a secret piece of news is divulged by a spy before the time is ripe, he must be put to death *together with the man to whom the secret was told.*”

Emphasis added. In this report we want to suggest a more nuanced, though not necessarily more effective, approach to the problem of maintaining information security in a warfighting environment.

1.2 Today's Context Is Important for This Study

Tension between those who produce secret information and those who consume it is not new. Content producers uniquely understand the fragility of their sources and methods in all detail. Their perspective is often a long-term one. The value of today's information from a particular source or method may be insignificant in comparison to the value of the future information that the source or method may yield – at almost any rate of discounting to the present. Yet a single misuse of today's information may jeopardize that whole future.

Content consumers, especially warfighters, often can not afford the luxury of a long-term view. When there is information whose use today can save U.S. lives (or enable warfighting objectives), *with certainty*, and when the possibility of losing a future information source is only *speculative or statistical*, then the warfighter's bias will be to use the information, now.

Both communities are attempting to make rational decisions about risk versus value. However, because they do not share the same risk model, or have an agreed-upon framework for measuring value, it is natural that they will often reach different conclusions.

Today, the historical tension between these two communities seems near the breaking point, threatening to tilt the previously stable (though not necessarily optimal) balance to one that favors the warfighter.

Four factors are destabilizing to the meta-stability of recent decades.

First, of course, is the end of the Cold War. Global nuclear strike against the U.S. by a peer power is no longer the fulcrum on which defense posture balances. But the end of a cold, symmetric threat environment turned out to demarcate the beginning of a hot, asymmetric threat environment. Warfighting is not a hypothetical future activity for the military; it is here and now, and at a tempo that seems to be increasing with time.

Second, and related, are the new challenges of homeland security. In the war against terrorism, our front-line troops are not all soldiers, sailors, fliers, and marines. They are also police, firefighters, medical first responders, and other civilian personnel. These are groups whose historical access to sources of national intelligence has been near zero; yet their need for real-time and analytical intelligence is now critical.

Third, is force transformation in general, and net-centric warfare in particular. Our defense posture is irrevocably committed to substituting information dominance for "heavy armor" (in all senses). Smaller units, lighter units, distributed forces, remote fire, precision strike: virtually every doctrinal idea under the "transformational" umbrella is enabled by the rapid flow of intelligence and other information. If this is to be the way that we fight (and it will be), then information producers must either learn to support this model or go out of business.

Fourth, is the notable change in the character of the professional military, especially the enlisted ranks. Soldiering is now a high-tech profession. We expect of the individual fighter a set of technical and judgmental capabilities, and a level of individual initiative, that is arguably unprecedented – and certainly unprecedented since the rise of mass conscript armies in the time of Napoleon. The capital cost of equipping a soldier today is probably two orders of magnitude greater (after correcting for inflation) than in World War II. Recruits today have grown up with video games, instantaneous web access, cell phones, instant messaging, CNN and Fox News, instantly downloadable maps and driving directions, and a host of other technologies that could translate directly into the combat environment. These soldiers have high expectations for warfighting technologies in general, and information technologies in particular. The consumer of intelligence is no longer an O4 "behind the green door." She is an E4 behind the (camo-) green door of a humvee – and it is *moving*.

The significance of the last three factors is amplified by the ever increasing pervasiveness of personal computing and the world-wide web.

1.3 What Problems Need Solving?

In the above context, it is easy to list the key problems that need to be solved.

- Information flow to the warfighter is perceived by many to be – and we concur in this judgment – excessively constricted.

We need new technologies for acquiring, merging, and delivering the information faster and in a more useable form; and we also need new information security constructs so that the full value of the information can be realized by delivering it to the broadest set of users consistent with its prudent protection.

We must also change a culture in which the logically separable roles of “content producer” and “content protector” have become completely entangled, and in which “knowledge is power” is too frequently mutated to “*withheld* knowledge is power.”

- There is no presently accepted paradigm for providing intelligence and other classified information to distributed homeland security consumers.

This refers to both first responders (police and fire chiefs, etc.) and also to local government officials with other operational responsibilities (mayors, city managers, power and water officials, etc.).

- The gap between the implicit risk/benefit calculations of the producer and consumer communities is greater than it has ever been.

Users see an overly rigid, out of date, bureaucratic structure of information classification and originator-controlled distribution; and an individual

clearance process that is glacially slow, and under which large numbers of fighting men and women are, in practical terms, unclearable.

Information producers, and others charged with information protection, perceive nearly insurmountable new challenges to information security, fueled by burgeoning computer networks, new RF technologies, and newly capable foreign intelligence adversaries.

- The status of sensitive information outside of the present classification system is murkier than ever.

Certain work-arounds to the present system result in classes of information whose protection level is uncertain. "Sensitive but unclassified" data is increasingly defined by the eye of the beholder. Lacking in definition, it is correspondingly lacking in policies and procedures for protecting (or not protecting) it, and regarding how and by whom it is generated and used.

2 THE PRESENT SYSTEM OF CLASSIFIED INFORMATION

The present system for information protection defines “classified” information, and defines a “clearance” process by which individuals become trusted, that is, permitted to access classified information. By now (2004), all active defense professionals and military personnel in the U.S. have never known any other system. It therefore seems natural and immutable to us. This makes necessary change more difficult.

In later sections of this report we will show that the present system is unnatural, far from optimal, and that it ought to be radically changed. Therefore, we first want to review the conceptual framework of the current system with care.

2.1 The Present System Is Largely Unchanged Since the 1940s

2.1.1 Before 1940 There Was Only Informal Classification

Who can doubt that the protection of confidential information has been a necessary practice of all governments throughout all of human history? In the American colonies, predating the Declaration of Independence, the First Continental Congress passed this resolution on September 6, 1774:

“Resolved, that the doors be kept shut during the time of business, and that the members consider themselves under the strongest obligations of honor, to keep the proceedings secret, until the majority shall direct them to be made public.”

Military secrets, and criminal penalties for revealing them (e.g., espionage laws) similarly date from nearly the founding of the Republic. Of some relevance, however, is the fact that secrets were defined by their *prima facie* content, not by a system of "classification" in the modern sense – that is, an adjudication and marking of documents at the time of their creation.

The pinnacle of such "informal" classification systems must surely be the Espionage Act of 1917, whose first section is worth quoting in detail.

(a) Whoever, for the purpose of obtaining information respecting the national defence [sic.] with intent or reason to believe that the information to be obtained is to be used to the injury of the United States, or to the advantage of any foreign nation, goes upon, enters, flies over, or otherwise obtains information, concerning any vessel, aircraft, work of defence, navy yard, naval station, submarine base, coaling station, fort, battery, torpedo station, dockyard, canal, railroad, arsenal, camp, factory, mine, telegraph, telephone, wireless, or signal station, building, office, or other place connected with the national defence, owned or constructed, or in progress of construction by the United States or under the control of the United States, or of any of its officers or agents, or within the exclusive jurisdiction of the United States, or any place in which any vessel, aircraft, arms, munitions, or other materials or instruments for use in time of war are being made, prepared, repaired, or stored, under any contract or agreement with the United States, or with any person on behalf of the United States, or otherwise on behalf of the United States . . . ;

or (b) whoever for the purpose aforesaid, and with like intent or reason to believe, copies, takes, makes, or obtains, or attempts, or induces or aids another to copy, take, make, or obtain, any sketch, photograph, photographic negative, blue print, plan, map, model, instrument, appliance, document, writing or note of anything connected with the national defence;

or (c) whoever, for the purpose aforesaid, receives or obtains or agrees or attempts or induces or aids another to receive or obtain from any other person, or from any source whatever, any document, writing, code book, signal book, sketch, photograph, photographic negative, blue print, plan, map, model, instrument, appliance, or note, of anything connected with the national defence, knowing or having reason to believe, at the time he receives or obtains, or agrees or attempts or induces or aids another to receive or obtain it, that it has been or will be obtained, taken, made or disposed of by any person contrary to the provisions of this title;

or (d) whoever, lawfully or unlawfully having possession of, access to, control over, or being entrusted with any document, writing, code book, signal book, sketch, photograph, photographic negative, blue print, plan, map, model, instrument, appliance, or note relating to the national defence, willfully communicates or transmits or attempts to communicate or transmit the same and fails to deliver it on demand to the officer or employee of the United States entitled to receive it;

or (e) whoever, being entrusted with or having lawful possession or control of any document, writing, code book, signal book, sketch, photograph, photographic negative, blue print, plan, map, model, note, or information, relating to the national defence, through gross negligence permits the same to be removed from its proper place of custody or delivered to anyone in violation of his trust, or to be lost, stolen, abstracted, or destroyed, shall be punished by a fine of not more than \$10,000, or by imprisonment for not more than two years, or both.

What is interesting about this section is that it is *not* the part of the Act dealing with espionage *per se*. (That comes later.) Rather, its subsections (a) - (c) attempt to criminalize the receiving of (what we would now call) classified information *by anyone*, while its subsections (d) and (e) attempt to criminalize what we would now call misuse of classified information by cleared personnel. We say "attempt" because the Act triggered a complex set of judicial reviews (beyond our scope here), and by today's standards would seem especially unlikely to be upheld as constitutional.

The system of explicit classification that we have today, which dates from 1940, should be viewed as having had a goal of "classifying" more explicitly the information that should be subject to access control, and providing civil and criminal penalties for misuse of that information.

2.1.2 The Present System Was Established in 1940 and Has Changed Little

This section is based on information from the Office of Declassification's history of classification and declassification, dated July 22, 1996.

On March 22, 1940, President Roosevelt signed Executive Order 8381, which defined certain kinds of information that could be "classified" (a very broad set), and established three levels of information: Restricted, Confidential, and Secret. Also specified were "classification authorities" (persons who could classify materials); these were also very broad, essentially all military and civilian government employees. This Executive Order governed classification throughout World War II.

President Truman, in EO 10104, signed February 1, 1950, limiting classification authority to the new Department of Defense, and adding Top Secret to the previous three levels of classification. In EO 10290, issued September 24, 1951, classification authority was expanded to all executive agencies, military and non-military. This order first provided for downgrading and declassifying, either automatically or upon review.

From this point on, most changes have been either procedural tweaks, or else political rebalancing in accord with an administration's liberal or conservative bent.

The governing order from 1953 to 1972 was President Eisenhower's EO 10501, which reduced the number of agencies with classification authority, and which eliminated Restricted as a classification level. Here automatic declassification was first tied to a date or event specified by the original classifier. President Kennedy's EO 10964, in 1961, slightly modified the rules for automatic declassification, and added administrative sanctions against any individual who knowingly revealed classified information.

President Nixon's EO 11652, in 1972, responded to the Supreme Court's decision in the case of the Pentagon Papers. Certain types of information were prohibited to be classified. Paragraph marking was introduced. Some review provisions were made mandatory.

President Carter's EO 12065, in 1978, changed the definition for Confidential to require "identifiable" damage rather than just "damage." Also, it established a balance test to determine if public interest outweighed possible damage to national security. A policy of less restrictive classification in borderline cases was enunciated.

President Reagan's EO 12356, in 1982, effectively repealed the liberalizing provisions of EO 12065.

President Clinton's EO 12958, in 1995, in substance reinstated the Carter Administration provisions.

No administration since Roosevelt has made significant changes in the overall system for information protection: information is still "classified"; individuals are still "cleared."

2.2 In Principle, Risk Level Governs Classification

There are at present seven specified categories of information that can

be classified by an authorized original classifier.

- military plans, weapons systems, or operations
- foreign government information
- intelligence activities, sources and methods, cryptography
- foreign relations, foreign activities in U.S., confidential sources
- scientific, technological, or economic matters relating to national security
- programs safeguarding nuclear materials or facilities
- vulnerabilities or capabilities of systems, installations, projects or plans relating to national security

In practice, individual agencies and programs have phone book-sized classification guidelines that explicate and expand these categories within an agency's mission space.

It is required that the original classifier be able to identify or describe what damage could occur if the information were to be released. However, there is no requirement that this be recorded on a decision-by-decision basis, other than by the original classifier's decision (again based on detailed guidelines) as to whether to classify as:

- Confidential (implying "damage"),
- Secret (implying "serious damage"), or
- Top Secret (implying "exceptionally grave damage")

Certain categories of information do not fall within the classes of information that can be classified, but can be protected from public disclosure

under one of the nine statutory exemptions within the Freedom of Information Act (FOIA). Such "sensitive but unclassified" information can be marked For Official Use Only (FOUO). Its unauthorized release can result in civil penalties against individuals of up to \$500,000. However, agencies differ widely in their application of the FOIA exemptions and in their training of personnel in this subject. The protection of FOUO material is therefore much more variable than is the protection of classified material.

Some agencies have from time to time (especially post-9/11) attempted to establish categories of informally protected information that go beyond the listed FOIA exemptions. It is likely that these attempts have no enforceable statutory or regulatory basis, and that information so "protected" could be obtained by any individual under FOIA.

2.3 In Practice, Distribution Channels Are a Dominant Factor

What happens in practice is, of course, rather more complex than the previous section indicates, both because of additional statutory and regulatory provisions, and because longstanding institutions develop "cultures" that dictate behavior.

Much information, especially intelligence information, is "owned" by its originator from cradle to grave, i.e., not just for the purposes of an original classification decision.

Some, but not all, information producers control stovepiped distribution channels, from which the associated product may not be removed. Some, but not all, information producers can, by various mechanisms, create compartments within distribution channels, and can limit access to those compartments at the granular level of specific individuals.

Individual clearances are indeed keyed to allowed levels of information

access (C, S, TS), and these clearances are indeed based on background investigations, in principle reinvestigated every five years. But in practice, a user's access in a given context frequently depends less on their clearance level than on the channels and compartments to which they have been admitted. Admission to these channels or compartments, while justified in principle by need-to-know, can in practice often be an arbitrary decision made by those who own the channel or compartment, or a decision limited by the number of "billets" in the channel or compartment.

We do not doubt that need-to-know does inform most decisions about access. However, the manifest pressure for horizontal integration (resulting in this study, for example) clearly shows that not every user with need-to-know, as perceived by the user, is being served.

2.4 Computers and Networks Present Enormous Challenges

Computers and networks now dominate information processing and distribution, but the present system was certainly not designed with these in mind. Cyber information thus presents enormous new challenges to information security. These include

- Volume of information. A 1-gigabyte memory stick that can attach to one's keychain can hold about 500,000 text pages. At this level the quantitative difference becomes a qualitative difference.
- Rapidity of information distribution. "Oops, clicked the wrong icon" can have instantaneous and irreversible repercussions.
- "The" Internet, that is, the concept of a universal network. The extraordinary power of the Internet in large part derives from its unity. Even when we build air-gapped, separate networks, we use totally interoperable protocols and hardware. If the air gap is inadvertently or

maliciously bridged, the IP packets can go anywhere.

- Ease of duplication of information. The "insider threat," which includes both individuals and surreptitious malicious software, is hugely magnified. Also, on the Internet, the number of copies of a juicy document can grow geometrically. (A good example is the Church of Scientology's futile attempts at recovering its "stolen" intellectual property.)
- Nearly invisible information exfiltration. Not only is data exfiltration by an attacker enormously easier today than ever in the past, but it is also more invisible. It seems quaint to recall the days when briefcases were routinely opened for inspection when exiting certain classified facilities, as much to prevent accidents as intentionally bad behavior. Today, a classified facility's unclassified (hopefully separated from its classified) connection to the Internet may carry many gigabytes of unspectable information per day.
- Data mining. Computers allow the adversary not only to gather information, but also to actively mine it. A computer is much more than a passive repository of information.
- High noise level. Hackers (on the outside) and computer glitches and malfunctions (on the inside) create a background of events against which actual attacks may not be detectable.
- Offense-Defense gap. Many well-publicized Red Team successes show that, with current levels of technology and practice, it is easier to attack computer systems than to defend them.

Computer security professionals in government have responded to these difficult challenges as best they can. There are government-wide and agency-specific procedures for certifying classified networks and the computers attached to them. The certification requirements are generally very conservative (allowing little acceptance of risk). There are large volumes of rules and recommendations for best practices.

However, these certification requirements, rules, and best practices are in aggregate very difficult to implement correctly, particularly in operational settings. The rules are often technologically out of date, either requiring out of date (and often unobtainable) technology approaches, or else blind to new technologies and their associated threats.

The operator's response is inevitable and predictable: the rule books are not being followed. In the best cases, deviations are formally requested (and granted) by approved mechanisms. Often 180-day waivers are granted with the unspoken understanding between the operator and the regulator that "180 days" actually means "forever." In the worst cases, users become scofflaws, so that organizations responsible for network security do not even know the extent of the problems they face.

The current situation of out of date or operationally unimplementable rules, combined with widespread violation of those rules, is a bad place to be. It implies that risk is being managed (or, rather, not managed) in a highly inoptimal fashion. If, because of usability issues, relatively secure technologies are not used, then relatively risky behaviors are substituted. *We are thus in a situation where aggregate risk would be lowered by the use of less-secure - but more useable - technologies.*

2.5 There Is Increasing Evidence That the Present Construct Is Breaking Down

We have already mentioned several pieces of evidence suggesting that the present system is breaking down. Let us summarize them here:

1. Users are dissatisfied with the present system. They don't think it meets their needs.
2. The present system is inadequate in keeping up with technology, both from the user perspective, and from the perspective of those charged

with protecting information.

3. Users are electing to avoid entering the system entirely. For example, Predator imagery in Iraq is considered Unclassified, but is protected by an *ad hoc* system of operational practices.
4. The system is being distorted by operational needs. Underclassification of documents – often quietly justified as necessary for ease in transporting documents between meeting sites – is a well known practice. Some agencies avoid using encryption in applications where it would have a clear security benefit specifically because the use of encryption immediately requires NSA involvement. That agency is perceived (whether correctly or not) as overly conservative, technologically backward, and intensely bureaucratic.
5. Operational users have figured out that it is better to ask for forgiveness than for permission. We heard statements such as this many times from briefers: “If you ask for permission, you’ll be turned down, but if you just do it, and are meeting someone’s operational need, then if you’re ever found out, you’ll get a waiver.”

2.6 The Present System Has Some Good Constructs But Is Missing Others

Can we say anything good about the present system? Yes, actually.

2.6.1 Constructs Present in the Present System

Clearly the original intent of the present system was that it be risk-based. The triage of information into three categories (now C, S, and TS), and the requirement of a more rigorous clearance process for individuals to access the higher categories, is manifestly a system of risk management, as distinct from blind risk avoidance.

The foundational logical constructs of the present system are thus:

- A construct for describing the magnitude of damage incurred if a piece of information is compromised (C, S, or TS).
- A construct for rating the reliability of (i.e., degree of risk inherent in) an individual (uncleared vs. cleared to various levels).

The later growth of distribution channels is, in effect, a third foundational construct, although rather *ad hoc* in implementation

- A construct for distinguishing qualitatively different types of secrets.

There are other *ad hoc* constructs that have sometimes been both sensible and effective. For example, compartmentation is used to separate “fact-of” secrets (where the *existence* of a technology or source is a secret) from “fact” secrets (where the secret resides in the specific details). Stealth technology is a good example: During its long development phase, fact-of-stealth was successfully protected by a compartment that contained multiple “facts about stealth” subcompartments. When fact-of-stealth was revealed (by SECDEF Harold Brown on August 22, 1980 – an action that remains controversial to this day¹) the outer compartment was breached, but the inner compartments remained water-tight, as can be judged by the technical nonsense that appeared in the media at the time.

Another well developed, though initially *ad hoc*, construct embodied in the present system is the separation of source from product. Here, a successful example is Classic Wizard, a system characterized in the media as an advanced tactical ocean surveillance system. RADM Thomas Betterton has described his approaching a senior admiral to obtain the latter’s backing for improving Classic Wizard’s front-end collection system, which he had

¹A fascinating description of these events is at <http://www.airpower.maxwell.af.mil/airchronicles/apj/cunn.html>.

identified by its intelligence community name. "I don't need that," Betterton was told in no uncertain terms, "because I already have Classic Wizard."

(Sun Tzu's version of separating product from source is: "When all five kinds of spy are at work, none can discover the secret system. This is called 'divine manipulation of the threads'. It is the sovereign's most precious faculty.")

2.6.2 Missing Constructs in the Present System

One can readily identify logical constructs at about the same level of abstraction as those just described that are missing in the present system, and that may be useful (or essential) in any follow-on system.

- A construct for implementing responsiveness to operational necessity.

In the present system there is no way to turn up or down the knob that governs the tradeoff between security and operational needs. There is no way, in time of war or in a particular area of operations, to "moderately increase" all players' access to secret information. The paradigm of "information is classified, individuals are cleared" allows for no continuous variables.

Of course in times of stress some technical and administrative barriers that protect secrets are just dismantled completely. This is a rational response to a sudden need for more flexibility, but it is not well calibrated; there are too few steps between highly secret and totally open.

This limitation is particularly troublesome in today's multilateral and coalitional world, where we must deal with new partners in an agile, yet controlled, way.

- A construct for risk based on the type of access transaction.

The present system lacks any means of recognizing that not all access transactions present equal risk, or for measuring such risk quantitatively. For example, it is obvious that the one-time display of a classified document on a (secure) computer terminal to a (cleared) individual – which we can call “soft access” – is inherently less risky than providing that same individual with a paper copy of the same document – “hard access.” Indeed, risk in the soft access case is smaller for multiple reasons, including: 1. If the individual is untrustworthy, his ability to convey the secrets to a third party is much smaller. 2. If the individual is trustworthy but careless, the soft access method provides less opportunity for misplacing a secret document. 3. Even if the individual is trustworthy and careful, the time-persistence of a paper document (ordinarily months or years), and its unencrypted format, provides many more opportunities for physical compromise by others.

- A construct for incentivizing the distribution and utilization of information.

The present system conflates two logically distinct roles: content producer and content protector. While the producer doubtless has a keen interest in protecting information, that person is (or should be) also interested in ensuring that the information is used. The present system greatly skews the outcome of this intrinsic conflict of interest by explicitly tasking the content producer to also be the first-line protector of the information. In fact, this is the *only* rationale for producer (originator) ownership of classified information, a key principle of the current system.

- A construct for auditing.

We audit (sometimes) with the purpose of maintaining physical inventory control of classified items. We record individual classified accesses (sometimes, and for limited kinds of classified information). But very rarely do we systematically audit records of individual classified accesses. There are multiple reasons for this: 1. Too hard. 2. Too incomplete. 3. What would we

do with the results anyway? 4. No standard formats. 5. No infrastructure for aggregating audit results. 6. Only look when there is a problem that is already known.

- Appropriate constructs for non-documentary information.

Non-documentary information includes things like "sessions," which may be audio, video, workgroup software, or any combination; and services, including database searches, data mining, and so forth. In the present system, we treat these as being classified at the highest level as the information that they *could* contain. But this makes some transactions difficult. For example, one might want to allow Secret-level access to the Secret information in a Top Secret database – very difficult today.

3 TOWARDS A NEW SYSTEM FOR INFORMATION PROTECTION

3.1 A New System Should Satisfy Some Basic Criteria

Drawing on the previous Section 2.6, we can set forth some basic criteria that any new system should satisfy.

1. It should meet operational needs, both today and in the foreseeable future.
2. It should be risk-based, and measurably so.
3. It should be agile and extensible, accommodating variable risk acceptance and shifting coalition partners.
4. It should enable incentives for information sharing.
5. It should be simple enough to be explainable in common-sense terms.
6. It should be objective, not subjective.
7. Its decisions should be auditable.
8. It should recognize qualitative distinctions between different kinds of secrets (e.g., tactical vs. strategic)
9. It should be capable of being implemented incrementally.
10. It should be capable of balancing risk to technical attack against human risk so as to achieve lower risk overall.

3.2 The STU III Is a Good Example of Accepting Risk

The last item in the above list of desiderata merits additional discussion, which we offer in the form of a success story. The STU III secure telephone, first introduced in 1987, was enabled not only by technical advances in microelectronics and cryptography, but also by important changes in the willingness to accept risks to technical attack. The predecessors to the STU III, the original STU (1970) and the STU II (1975), were each the size of a small refrigerator. The cabinet of each was something like a steel safe, with a regulation combination lock securing it against unauthorized entry. The STU and STU II were handled as classified cryptographic equipment. They were shipped by classified channels and required the continuous protection accorded to classified cryptographic materials. At its peak, the STU II had about 10,000 subscribers.

By contrast, the enormously successful STU III (and its successor, the Fortezza-based STE) can be purchased by authorized users through unclassified commercial channels. It can be shipped using ordinary commercial carriers. Once activated, it can be left out in any room where classified discussion is permitted. Part of the key for a STU III is stored in its "crypto ignition key" or CIK. The rules for storing the CIK are as follows

"When the CIK and the STU-III are kept in the same room, the CIK must be protected at the highest classification level of the information that the STU-III is authorized to transmit.

When not kept in the same room as the STU-III, however, the CIK may be protected as you would a high-value item of personal property, such as a credit card. It may be stored in a locked cabinet or desk. It may also be kept in the personal possession of the authorized holder."

How was this enormous increase in utility brought about? Without doubt, the largest factor was a new set of cryptographic protocols, which made compromise of the information in a single STU III much less damaging than compromise of a STU II. However the STU III is not immune to the risk of non-cryptographic technical attack. For example, a STU III in

transit could have a malicious implant installed, one that exfiltrates classified discussion through a channel other than the secure cryptographic one.

It was a wise decision by NSA policy makers that the residual risks of the STU III to technical attack – risks that were not present with the STU II – were smaller (we would say vastly smaller) than the risk reduction that would be realized by attracting a larger user population with a more useable product. Today, several hundred thousand STU III phones are in use. Billions of phone conversations that might otherwise have taken place on open lines, “talking around” classified topics, or not take place at all (with loss of program effectiveness) have been attracted into the STU III infrastructure.

Several present and former NSA officials have told us that, in today’s climate, the STU III would never be authorized. The standard for technical perfection, and the aversion to accepting risks of technical attack, are today so high that they prevent the implementation of better end-to-end lower risk solutions. This is truly a case where the perfect is the enemy of the good. A future security system should ensure that this skew does not continue.

3.3 IAD’s “45 Day Study” Is Useful But Not Radical Enough

In response to urgent tasking, NSA/IAD chartered an Assured Sharing Tiger Team, whose report, “Access Control Concepts for Assured Sharing” was delivered on May 26, 2004. JASON was asked to comment on this report.

We find that the 45 Day Study is a useful contribution to the horizontal integration discussion, but that it does not in itself provide a roadmap for future action.

In brief summary, the 45 Day Study proposes to augment the current information security construct with a process for approved exceptions, whereby

– in response to urgent operational need – access to classified or compartmented material could be granted to individuals without the appropriate clearances and distribution channel approvals.

The Study details at some length how this “Risk Adaptable Access Control” (RAdAC) might function. The Study document is a good source for taxonomic lists of the various pieces of infrastructure, characteristics of information, attributes of people, and so forth, that should or could be brought to bear on access decisions.

However, the proposed approach addresses only some of the desiderata of Section 3.1, above, and falls short in a number of them:

- It is questionable whether the complex new layers of process that it introduces (the exception process) will meet operational needs.
- Although it furnishes infrastructure for the granting of exceptions, and potentially provides the exception-granter with data relevant to risk (e.g., personnel attributes), it does not itself contain a risk model. That is, it gives no guidance to the exception-granter as to what that person should do or not do. This would result, we think, in highly variable decision-making.
- Although it enables information sharing, it does not incentivize it.
- It is questionable whether the process is simple enough to be explainable in common sense terms.
- Because actual decisions are simply pushed to the exception-granter, they are not particularly objective, and they are auditable only as to action, not as to reason.
- It appears to envisage “perfect” physical infrastructure; that is, it does not have a mechanism for acceptance of risk in less-than-perfect instantiations.

In conclusion, we do not think that the 45 Day Study in itself is a sufficiently broad, or sufficiently far-reaching, approach.

3.4 Three Guiding Principles for Any New System

We recommend that any new system focus on *risk*, and not waver from that focus. We propose that three principles, outlined here, be applied. Below, in Section 4, we will give a specific proposal for a new system; but the principles given here are more general than that specific proposal.

1. *Measure risk.* “If you can’t measure it, you can’t manage it.”² When risk factors can’t be measured directly, they can often be plausibly estimated (“guessed”); and subsequent experience can then be used to derive increasingly better estimates. This can be formalized in various ways, for example as an updatable Bayesian belief net.

2. Establish an *acceptable risk level*. As a nation we can afford to lose X secret and Y top secret documents per year. We can afford a Z probability that a particular technical capability or HUMINT source is compromised. If we set X, Y, Z, \dots all to exactly zero, then all operations stop, because all operations entail some nonzero risk. What, then, are acceptable ranges of values for X, Y, Z and a long list of similar parameters? Even better, roughly what values would *optimize our operational effectiveness in the long run*, when today’s security breaches show up not simply as harmful in the abstract, but as actually imperiling future operational effectiveness?

3. Ensure that information is distributed *all the way up to* the acceptable risk level. This is a very fundamental point. We have been living with systems that try to minimize risk. That is the wrong metric! We actually want to *maximize information flow*, subject to the overall *constraint* of

²The quote is due to Peter Drucker, but is often incorrectly credited to Demming. Demming actually insisted that good managers must strive to make good decisions even in the case of very incomplete data, which likely is the case in many areas of information security.

not exceeding the acceptable risk levels set according to principle number 2, above. This means that instead of minimizing risk, we actually want to ensure that it is increased to its tolerable maximum (but no higher).

These principles are not a roadmap. They tell us what we need to achieve, but not how. In the next sections we begin to suggest how one might design a new architecture for information security based on these principles.

4 SPECIFIC PROPOSAL FOR A NEW PARADIGM

It is not fair for us to shoot arrows at other people's efforts unless we are willing to let ourselves also be a target! In that spirit, we use this section to propose the broad outlines of a complete, new architecture for information security. We will start with guiding principles of Section 3.4, and then suggest how they might be implemented in a specific "point design." We fully expect that much of what we suggest can be proved unworkable, or no better than some different approach. But by putting forth an ambitious proposal, we hope to make a contribution toward advancing the discussion in a concrete way.

4.1 We Propose a Three-Phase Plan

There are three phases in the plan that we will outline: a preparatory phase (Phase 0) during which necessary prerequisite infrastructure is built, some directed research is done, and pilot programs are initiated; a first operational phase (Phase 1), in which the three guiding principles are actually implemented, but the present construct of producer (originator) control of distribution is not completely overturned; and a more ambitious second operational phase (Phase 2) in which originator control is replaced by something that one might call a "market economy," satisfying the three guiding principles and optimizing the use of information under those principles.

4.2 Preparatory Phase (Phase 0)

We must accomplish two goals during the preparatory phase. First, we must develop a risk model based on measurement (or estimate) of transactional risk. Second, we must develop the infrastructure (physical and pol-

icy) necessary to support the transactions that are contemplated by the risk model. We elaborate on these goals in sections 4.2.1 and 4.2.2, respectively.

Additionally, some pilot programs can be fielded during the preparatory phase, both to meet immediate needs, and to establish confidence in the concepts and approaches that will be implemented in later phases. This is discussed in section 4.2.3.

4.2.1 Developing a Risk Model

We advocate moving from a system based on (what amounts to) lifetime trust in cleared individuals to a system based on accounting for the risk of individual transactions. By "transaction" we mean events like:

- an individual accesses a classified document on a computer screen,
- a classified briefing is broadcast to multiple individuals,
- a classified document is printed in hardcopy and stored in an approved safe,
- an uncleared soldier at a checkpoint in Iraq checks a cell phone number against a classified data base (not possible under the present system),

and many other similar events.

Needed is a model that quantifies the risk of different transactions, both relative risk ("hardcopy is X times more likely to result in information loss than softcopy"), and absolute risk ("there is a 10^{-Y} chance per exposure that this secret will end up published in Aviation Week").

A transactional risk model should include factors that relate to the individual(s) involved in the transaction, for example:

- U.S. citizen vs. foreign national,

- personal history (rank, length of service, etc.),
- clearance level, and surrogates such as credit rating, arrest history, etc.,
- quality of the security brief-in given to the individual,³
- the individual's history of access to classified information (someone with very broad accesses may present an intrinsically higher risk than someone with narrow previous accesses, but this could be offset by a factor like that in the preceding bullet),
- ability to sanction the individual (that is, credible fear of punishment).

It should also include factors that relate directly to the type of transaction, for example:

- display only vs. hardcopy,
- one time vs. multiple access,
- volatile vs. aggregatable or archivable to local storage,
- access is auditable vs. non-auditable,
- redistribution is infeasible vs. feasible.

Finally, there are situational factors surrounding the transaction that are not directly attributable to either the individual or the means of access. A few examples are

- how controlled is the environment? (is it a SCIF, a battlefield HQ, or a more vulnerable location?),

³Several JASONS, with security clearances going back many decades, have noticed a dramatic decrease in the "quality" of security brief-ins. Brief-ins were once given personally, by a figure of some authority, and with the flavor of a personal transference of responsibility and trust. Now, more often than not, a security functionary says "sign here," or at best "sign here that you've watched the video." We think that the quality of early brief-ins can directly affect how seriously the individual takes his or her security responsibilities – for a lifetime.

- how directly can we link the nature of the access request to the operational environment? (transactions that, on the basis of objective criteria, appear "out of routine" may be assigned a higher risk than transactions that are evidently linked to an individual's or unit's assigned work),
- how secure is the technology utilized? What level of encryption is employed? What is the risk that the user's equipment is compromised?

We think that a directed research program, with milestones over a period of about two years, could yield a serviceable transactional risk model that integrates factors such as those listed above into a single algorithmic framework – a "formula." The model need not be perfect to be useful, but it must be objectively based on available data from a variety of sources. Such sources include systematic study of known espionage cases, statistics on security infractions and their consequences across the classified enterprise, expert opinion from CI professionals and experts on human reliability in high consequence industrial settings, and a variety of other sources.

The results of this directed research would enumerate failure modes of interest. On the human side, it would distinguish among carelessness, accidental disclosure, malicious disclosure, principled disclosure (the whistleblower or memoir writer), and so forth. On the technical side, it would distinguish among various kinds and levels of technical attack, for example tasked collection by a peer foreign intelligence service, serendipitous collection, hacking by tasked or untasked perpetrators, and so forth.

For each failure mode, we want to ask what empirical data is available, and how that data might be augmented by additional research activities during the Phase 0 preparatory phase.

Research should also address the question of how various risk factors, once identified, should be mathematically combined. It is unlikely that the answer is just a set of linear weights. Rather, certain combinations of factors

may enter into the risk in nonlinear ways, either reinforcing one another, or else possibly mitigating each other.

One might also ask whether some profiles present significantly different risks for different kinds of information, e.g., risk of disclosing short term secrets versus risk of disclosing long term strategic secrets. If so, this could usefully be incorporated into the overall risk model.

A balance must be struck between the risk model's complexity and its usability. Its development should not grow to become a monumental effort at classifying all transactional risks atomistically in all detail. Many cases will be duplicates or minor variants on a basic theme. In many situations, it might be good enough to aggregate risks at the unit level, and for a certain period of time. For example "use by a single Army company of Secret-level imagery of its operating area for 90 days" might be a single risk transaction. Indeed, this kind of aggregation might be a good way to start, with finer grain accounting introduced over time as the the supporting infrastructure becomes more capable and user-friendly.

The present system expends enormous resources on routine security investigations. An important function of the risk model will be to target those resources more effectively. In Section 4.4.5 we will give an example that shows how this might happen.

4.2.2 Developing Necessary Infrastructure

Also during Phase 0, significant new infrastructure must be developed in support of measuring and accounting for risk on a transaction by transaction basis. We comment here only briefly on a few of the most important prerequisites.

1. Protocol stack for moving and displaying protected information. By analogy with other protocol stacks in computer science (e.g., the 7 layer OSI

model of the Internet), what are the appropriate layered concepts for moving and displaying protected information across a wide variety of platforms, with support for both cryptographic infrastructure, and for the accountability of quantified risk? This is not a "green field" question. An enormous amount of work, both theoretical and practical, has been done on "digital rights management systems" (DRMSs). Do any of these systems form a basis on which we might build? Are there protocol layers in common? What new layers do we need?

2. Technologies for the controlled display and restricted use of information. Here also a large body of work already exists in the DRMS community (although the value may be more in the conceptual frameworks than in specific technologies).

3. User authentication technologies. COTS is progressing rapidly in this arena, particularly with regard to biometrics. What are the best COTS or GOTS technologies for our intended use? Note the importance of having a protocol stack that can accept many different kinds of authentication, and make use of the specifics of the authentication in assigning risk to a transaction (i.e., we must have appropriate communication between layers).

4. Technologies for limiting risk in less-controlled environments. The canonical case might be the ruggedized computer in the battlefield vehicle in theater. Tamper-proof chip technology should continue to be explored, noting that it is useless if it is always behind in "chasing COTS," with rapid commercial product cycles, to maintain compatibility. Some *interface* standards may have slower product cycles, however, and present opportunities for introducing tamper-proof security into those interfaces. Hard disk drive encryption may be one such example. Another example might be secure cards that are bus-mastering on a less-secure PC, and which would have the ability to monitor behaviors by that PC (akin to what network intrusion detection systems do on networks today).

5. Risk models for quantifying the acceptance of technical risk onto technical systems. For example, one has a sense that Security Enhanced Linux would be more secure if run off ROM than off a hard drive (see Section 4.2.4, below). But how much more secure? Can we quantify this in either absolute or relative terms, based on empirical data about the difficulty of past attacks (either against us or by us)? This is a difficult area, but that doesn't mean that we can afford to ignore it. We need to be able to make tradeoffs between human risks and technical risks (e.g., as in the example of the STU III in Section 3.2 above), and it is desirable to make such tradeoffs with some quantitative basis.

4.2.3 "Enclaves" or Communities of Interest

Phase 0 is not just a research stage. Interim and pilot programs can also be undertaken, both to reap immediate operational benefits, and to gain experience at some concepts that will become more widespread in Phase 1.

As one specific pilot program, we recommend that procedures be defined and implemented for creating temporary "enclaves" or "communities of interest." Within an enclave, more flexible access to classified information would be permitted, according to operational necessity.

In today's context, there is a need for a process to define "tactical compartments." An example of such a tactical compartment might be one applying to coalition ground forces in Iraq; or, more restrictively, geographically separate compartments for each of the major Iraqi cities and its surrounding area. Membership in such a compartment should be precisely defined, with an auditable roster maintained. However – and this is important, and different from any past practice – it would not be limited to personnel who hold security clearances under the present system.

It would be a matter of judgment what kinds of now-classified information to make available within a tactical compartment. All Secret-level

intelligence products, especially imagery, are obvious candidates. Certain SIGINT products with high tactical utility to the "soldier in the humvee" or "soldier at a checkpoint" also deserve consideration. There should be highly streamlined procedures for getting perishable information – useful now but not later – into tactical compartments, both for the benefit of its utility, and because the future security risk is small.

In some cases, intelligence products might be tailored to the tactical compartment. An example briefed to us was the concept of giving U.S. personnel at checkpoints the ability to check an Iraqi civilian's cell phone number against a classified watch list. In this case, the tactical compartment would not contain the watch list in aggregate – and certainly not information relating to its sources and methods – but just the *service* of checking the watch list in an automated (and perhaps rate-limited or "metered") manner. This example illustrates the general feature that *services* can be provided at a lower level of protection than the data from which they derive; it is a logical next step beyond the existing doctrine of separating source from product.

Documents within a tactical compartment should be appropriately marked as such, and there should be administrative penalties for mishandling such documents. The goal is certainly *not* to have a free-for-all release of classified information. Rather, the goal is to define, responsibly, a user group and set of products releasable to that group that is cross-cutting to the present system of classification and clearances.

Initially, the enclave should be viewed as a pilot project, an experimental way of doing business, as we prepare to make much greater changes in the present system. Nevertheless, legislative action is almost certainly required to enable even the relatively modest proposal for an enclave experiment that we recommend. The reason is that under existing law it generally is a criminal act to provide classified information to uncleared individuals, which is exactly what we may want to do in an enclave. We recommend that enabling legislation be pursued in the specific framework of legislation in support of Operations Iraqi Freedom and Enduring Freedom, where there might be more

than ordinary Congressional willingness to consider changes in ways of doing business.

Because the enclave is a pilot project, data – ground truth – should be gathered about how it comes to be used. In particular the data should support an ongoing risk assessment. How many classified images are shown to foreign nationals? Are the individuals coalition partners, who are likely to make information known to their home governments? Or, are they U.S. permanent residents without known ties to a foreign government? What U.S. sources and methods are likely to have become known to which foreign governments in the course of the enclave's operation? Is this loss tolerable in view of the operational benefits gained? (That is, did the enclave “make a profit”?) These kinds of assessments will be valuable not only in their own right, but also for developing processes for managing risk that will become central to later phases of implementation of the new security system.

4.2.4 NetTop Is an Important Technology

NetTop, an NSA project, is a technology for virtualizing machines and networks at different security levels onto a single hardware platform. The NetTop user has a single desktop or laptop machine, and has the ability to open windows on “virtual machines” at different security levels. Each virtual machine has its own operating system (e.g., Windows XP or Linux), and runs as if it were on a physically separate machine, connected only to the appropriate level of classified network.

NetTop does *not* instantiate so-called Multi-Level Security (MLS). MLS implies a single, unified, desktop with multiple levels of information on it, and a means of somehow enforcing the rules for upgrading and downgrading that information. Rather, the NetTop model is much more modest; but it nevertheless holds the prospect of significant immediate utility for the operational user.

NetTop instantiates and extends the concept of Multiple Signal Levels (MSL), rather than Multi-Level Security (MLS). MSL is the concept, fully accepted by the information security community, that different security levels can share the same physical hardware infrastructure, as long as high-grade cryptographic techniques are used to keep the differing levels completely separated, even when they are in the "same wire." Indeed, MSL is today's way of doing business in long-haul network transport. The encrypted text output of an NSA-approved crypto device is (generally) *unclassified*. It can, as a rule, be mixed with ordinary unclassified traffic, and sent over unclassified military or civilian channels.⁴

We are emphasizing this point only because there still seems to be some confusion about it in some tactical communities. We have heard an anonymous (apocryphal?) commander quoted as saying that he couldn't afford to bring five separate networks into his HQ for the five different levels of classification that had to be kept completely separate. If his meaning was "five separate pipes back to CONUS," then this view is simply uninformed: encryption products exist today to combine any number of security levels *while in transit*.

It is possible, however, that the apocryphal commander was not so badly informed, and that what he was referring to was the difficulty of having five separate sets of workstations inside his HQ; and five separate sets of wiring in the HQ between the NSA crypto boxes and these workstations. That is exactly the problem that NetTop can solve: NetTop pushes the MSL solution out of the wiring and onto the desktop. In a NetTop solution there are still (in this hypothetical example) five separate classification levels, all kept separate; but there is only *one* physical infrastructure, and *one* machine on the desktop. The apocryphal commander's concerns are thus addressed.

The technology underlying NetTop relies on two existing COTS or GOTS products. Closest to the physical machine, NetTop utilizes Security

⁴In rare cases, the threat of traffic analysis must be considered and mitigated. That discussion is beyond our scope here.

Enhanced (SE) Linux as its base operating system. SE Linux has the almost-unique distinction of being both NSA-sponsored and open source. That is a very powerful combination: the incentive structure of the open source movement motivates (probably) hundreds of very high-skilled programmers to be constantly scrutinizing SE Linux source code for security holes. The idea of publicly "showing up" NSA by finding and publicizing a hole has been a great motivator. Of course, holes found in this way have been immediately fixed. Meanwhile, NSA is able to apply its own expertise to the code without necessarily publicizing the full rationale for its changes. In NetTop, SE Linux runs in a particularly secure mode, without an IP stack (the point of entry for remotely hacking ordinary computers).

The second major piece of technology underlying NetTop is VMWare, a commercial product (now a division of EMC Corporation). VMWare runs on SE Linux, where it creates virtual machines and keeps them separate; not merely separate, but basically invisible to each other. NSA has performed a wide variety of tests and strong assurance activities on the version of VMWare that runs in NetTop (beyond our scope and classification level to discuss here).

Currently NetTop is certified by NSA for use in only one, restricted, scenario: it can be used to host separate Secret and Top Secret virtual machines, provided that the physical hardware is authorized for Top Secret level. That is, it allows hosting a less-classified network on a more-classified platform.

We recommend that NSA immediately consider additionally certifying the use of NetTop (in a configuration defined by them) to host separate Unclassified and Secret virtual machines on Secret-level platforms. (It is likely that the certified configuration would include at least one virtual machine, not controlled by the user, that performs required firewall and security monitoring functions on the unclassified connection.) We judge that this additional certification, with appropriate technical safeguards, carries negligible additional risk.

Perhaps more controversially, however, we also recommend that NSA consider certain scenarios in which higher-level virtual machines could be hosted on lower-level platforms, for example hosting the soft-copy display of tactical Secret level imagery on Unclassified (but protected) machines within a specified military area of operations. This would be a very timely example of *additional risk acceptance* onto technical systems (cf. the example of the STU III in Section 3.2 above) in the interests of large payoffs in both operational utility, and in avoiding the even larger risks of human workarounds and bad behaviors.

Various risk mitigations could be enforced in such a scenario. The Unclassified machine (running SE Linux) might have no disk drive, and load SE Linux from an unalterable ROM. It might be physically secured by an appropriate anti-tamper technology (fully electronic? also with physical tags and seals?), and with periodic inspections required. All of its non-ROM storage (main memory and RAM-disks) would be volatile, and erase at power-off. No Secret material would ever remain on the unpowered machine. The Secret level virtual machine might itself be SE Linux, running in a stripped-down mode and able to run only one, or a small number of, user applications for accessing the approved Secret level resources, and no others. A high degree of logging might be enabled, and logs could be moved across the Secret communications channel to an external secure facility in real or near-real time for automated checking and monitoring, with the ability to disable and lock-out a machine with a suspicious profile.

While NetTop is certainly not the answer to all problems, we think that it could rapidly deliver some very valuable interim solutions. We urge that it be given greater institutional prioritization at NSA, and that the Net-Top project be specifically tasked with developing moderate *risk acceptance* solutions that go beyond current practices, for limited rapid deployment.

4.3 Risk is Tokenized in Phase 1

The defining feature of Phase 1, the first operational phase for which the guiding principles (Section 3.4) are fully implemented, is that risk is “tokenized.” Since “tokenization” may be an unfamiliar concept, we will first explain what we mean by it in some detail.

4.3.1 What is a Token?

A token is something that has exchange value. That is, the holder of a token can trade it for something that they want. In the present application, that “something” is access to a specific piece of classified information by means of a specific “transaction” (see Section 4.2.1 for examples of different kinds of transactions). Some transactions may cost more tokens than others, just as certain goods in commerce may be offered at higher prices. So we can think of tokens as a currency, with the possibility of exchanges accounted for with any desired degree of precision (e.g., 3.4675 tokens).

Like currency, tokens must be instantiated in a way that makes them impossible to counterfeit, impossible to duplicate, and easy to move. Currency has some other attributes (like anonymity) which are not necessary in our application. We may demand some attributes (like auditability) that ordinary currency lacks. The tokens that we discuss will only be instantiated on computer networks; no one will be minting actual physical coins or paper money. On a computer network it is straightforward to instantiate tokens in much the same manner as any of the various “cybercash” protocols now used in internet commerce.⁵ That means that authorized individuals or organizational units have “accounts” that hold their tokens; and there are cryptographically protected, and generally user transparent, protocols

⁵Distinct from so-called “digital cash” protocols. Cybercash protocols do not guarantee many of the things digital cash does, like anonymity, and they rely heavily on a trusted bank.

for paying out tokens in exchange for goods, or transferring them to another person's account. Technologies exist that enable tokens to be verified, moved, or exchanged even in the absence of continuous network connectivity.

4.3.2 How Are Tokens Denominated?

Just as the dollar was once "pegged" to be 1 ounce of silver, or 1/20 ounce of gold, our tokens are pegged to risk. To be specific at the risk of oversimplification, we might peg the value of a token in Phase 1 by the baseline:

1 token = risk associated with one-day, soft-copy-only access to one document by the average Secret-cleared individual.

It is an important point that tokens in Phase 1 represent only the *risk* side of the equation (associated with a probability of loss or compromise), not the *value* side (damage due to a loss). We have not yet tokenized the value of a document – that comes only in Phase 2. But there is much to be gained even by tokenizing risk only.

In Phase 1, we can apply the risk model that was developed in Phase 0 to assign a token "cost" to different kinds of transactions, depending on whether they are riskier or less risky than the baseline transaction that pegs the value of the token. For example, the risk model might generate numbers like:

- one-day, soft-copy access for a Top Secret-cleared individual: 0.2 tokens (i.e., an individual cleared to TS poses less risk than one cleared only to S);
- one-time, soft-copy access for an *uncleared* U.S. citizen, married, with a good credit rating, who signs a legally enforceable secrecy agreement: 10 tokens;

- *hard-copy* access by a Secret-cleared individual with administrative restrictions on further copying: 50 tokens;
- hard-copy access by a Secret-cleared *with permission to redistribute* to other cleared individuals: 750 tokens.

These are completely notional numbers; the risk model will provide numbers with a sound basis founded in empirical data and expert opinion.

With tokens pegged to risk, and with a risk model that evaluates the relative risk of different kinds of transactions (including all of the individual-linked and transaction-linked factors discussed in Section 4.2.1 above), it is possible for the producers of new information to make judgments regarding total acceptable risk *denominated in tokens*. For example, it may be an acceptable security exposure for the average Secret-level IMINT photograph to be viewed, in soft-copy, as many as 1000 times by Secret-cleared individuals. In that case, each document created would have 1000 tokens created along with it, because the value of the token is pegged at one transaction of this type. (In modern terms, the tokens are a kind of metadata globally associated with the document.) If those tokens are spent not for 1000 accesses by Secret-cleared individuals, but instead for 500 such accesses plus 50 accesses by uncleared U.S. citizens under the conditions listed above (cost 10 \times), then the same maximum risk is “expended.”

Note that the tokenized risk model both allows greater flexibility and drives good security behaviors. For example, it allows limited access to classified information by uncleared personnel when such access is so important that a token holder higher in the chain of command is willing to pay the price. At the same time, it encourages the use of (cheaper-in-tokens) soft copy display instead of riskier paper copies of classified documents.

In practice we would expect the decision about how many tokens to allocate to a document (how much risk to deem acceptable) to be governed by policies set down by the producer’s organization at the agency level, and

not subject to the whims of an individual. In the example above, tokens are created only when the document is created.

In practice, it may be better to refresh tokens annually, so as to encourage continuing use of classes of documents, with specified *annual* risks. When a document, or class of documents, becomes less sensitive with time, greater numbers of its tokens can be issued in succeeding years.

The numbers above are purely notional. In actual implementation one would start with numbers of tokens created that closely approximate current usage patterns, so that, at inception, access to classified materials is neither easier nor harder than under the present system. Subsequently the token currency could be slowly inflated or deflated according to national needs, and in the light of experience.

4.3.3 Phase 1 Tokens Are Not Fully Fungible

We have thus far skirted the issue of whether there is only one type of token, or many incommensurable types. The answer in Phase 1 is: many. That is how Phase 1 preserves the principle of originator control of information even as it tokenizes aggregate risk. (Originator control will not be preserved in Phase 2.)

Thus, in Phase 1, one should think of one's token bank account as holding, in effect, different kinds of foreign currency: marks, guilders, lire, zloty, etc. Less flamboyantly, let us designate different kinds of tokens by the letter T followed by a 6 digit number:

- T110000 designates the token for Secret GENSER documents,
- T120000 designates the token for Secret-releasable IMINT products,
- T130000 designates the token for Secret-releasable SIGINT products,

- T503000 designates the token for COMINT-channel-only Top Secret products,
- T714568 designates the token for a certain highly restricted compartment,

again, all notionally.

Example 1: Annually, a new supply of T120000 tokens – exchangeable for access to Secret-releasable IMINT imagery – would be created by the National Geospatial-Intelligence Agency (NGA). In aggregate, this supply would represent the tolerable risk associated with all use of the product associated with this token. If the risk model has done its job properly, then the aggregate risk will be at about the same tolerable level independent of the type of transactions that occur.

Example 2: Separately, the NSA creates an annual supply of T503000 tokens that are exchangeable for TS-level COMINT. The number of tokens of this type created for this restricted product is much smaller than the number created in Example 1. That is because all tokens are pegged to risk in the same way, but the tolerable risk for this set of products is much smaller.

4.3.4 How Are Tokens Distributed?

One might just as well ask how is money distributed by an agency within its Congressional allocation? Or how are FTEs allocated to projects? The answer in all these cases is that it is a *management* function to allocate resources, and that resources are allocated hierarchically through a management chain or chain of command. Risk is just another limited resource, and risk tokens will be allocated, along with other resources, by management chains.

Thus, in Example 1 in the preceding section (4.3.3), NGA would allocate all created tokens, annually (say), to its users at the highest management

level, notionally to the Secretary of Defense, Secretary of Homeland Security, Secretary of Energy, and so forth. How this allocation is made is a national level decision, based on national needs. Doing so allocates risk according to expected return (that is, priority or need), with the constraint of keeping total national risk constant.

Within a Department or Agency, tokens are similarly allocated to consumers, down to the level of the first-line manager. Note that this is a "command" or "push" economy, not yet a "market" or "pull" economy. In Phase 2 we will change this. However, even in this command economy, the tokenization of risk drives good security behaviors (utilization of "cheaper" transactions) and encourages managers to locally optimize their use of classified resources for maximum returns to their missions.

Incidentally, it is also a management function, at all levels in the management chain, to monitor whether tokens are being spent wisely in support of the mission. Just like any other resource, management will pull back and redistribute surplus tokens, limit year-to-year carry-forwards, and so on. It will be hard for an individual to "hoard" tokens beyond their actual need, because they are so easily tracked by management.

4.3.5 How Are Information Producers Incentivized?

The tokenization of risk, when combined with the requirement that risk tokens be distributed all the way up to the annual tolerable risk, enables an elegantly simple figure of merit for information producers: What fraction of their tokens distributed are actually *used*, that is, exchanged for access to their classified products? We can call this measure the "utilization fraction."

If the utilization fraction is much less than 100%, then the nation is not getting the full benefit of these products, as allowed by the tolerable risk of their loss. This reflects badly on the information producer, who should be

held accountable. This is the implementation of Guiding Principle 3 (Section 3.4).

If an agency's classified product is not being used, as measured by utilization fraction, then the agency needs to do one or more of the following three things:

1. Improve the product, so that it is used more by existing customers.
2. Find new customers, and reallocate risk tokens from customers who are spending less than their allocations, so that the new customers are able to access the product.
3. (Rare.) Demonstrate that risk is logically not a limiting factor for this product. That is, show that within the tolerable risk level the product is already serving *all* possible users, and that its use by those users justifies its cost.⁶
4. (Even rarer.) Allocate fewer tokens to the product – that is, less than its tolerable risk level – and use the surplus tokens, representing tolerable aggregate risk, for another product.

But here is what the producer is *not* allowed to do:

- Jack up the claimed sensitivity of the product so that fewer risk tokens are created for it (market dislocation!)

This illustrates why it is important that the tolerable risk for individual products be set (or adjudicated) by someone other than the information producer. Because Phase 1 has, by intent, not broken the chain of originator control, the basic weakness of originator control is still present in the form of the above market dislocation. However, this “bad behavior” is at least exposed

⁶One can imagine a kind of Federal Register for underutilized products looking for additional government customers who could get utility out of spending unused available risk.

to external scrutiny ("Why are you issuing so few tokens for product ABC, by comparison with another agency's products DEF?"), which is usually not the case today.

Note that, as a success measure, token utilization can be tracked all the way down to the individual activity or analyst. When the products of many individuals are separately identifiable, but they are lumped within a single token ("T130000"), we can evaluate the aggregate activity by its fractional utilization, and we can evaluate the contributions of individual participants by the actual numbers of tokens spent to access the products that they individually produce.

As part of Phase I many details have to be worked out. For instance, the way of allocating, and accounting for the use of, tokens should be kept as simple as possible. In particular, the analogy with money doesn't extend to a need for an analog to double-entry bookkeeping, nor to the other elaborate procedures for ensuring that taxpayer money is spent properly. In a similar spirit of simplicity, there might be a token schedule for presentations, rather than a charge for each person in the audience, and a charge for continued use by a small organizational unit, rather than a charge for each access by each person. This is especially necessary if these are not computer accesses.

4.3.6 Steps Toward an Efficient Market Economy

Incentivizing by utilization fraction also encourages agencies to aggregate their products into smaller numbers of broader token types, so that their really popular products can help support the real dogs. This is actually a good thing, because – on the user side – it makes progress towards a freer market with a more fungible currency.

Additional steps could be taken in Phase 1 to encourage a more efficient market and the increased optimization of resources that such a market can give.

For example, one might encourage a secondary market in tokens, say (to keep the fig leaf of originator control), among those authorized to hold a particular token in the first place. That is, if you are allocated even 1 token of type T714568 (in effect being allowed into that compartment), then you can acquire more tokens, if you need them, on the secondary market. And what do you give in exchange? Why tokens of another type, of course! The market will quickly establish exchange rates among tokens, based on their utility to customers. By noting these exchange rates, higher level IC managers will be able to prioritize resources among different IC producers.

Note that in such transactions *risk is conserved*. Because the tokens represent risk, and higher risk transactions cost more tokens at the time that the tokens actually are exchanged for information, it is risk-neutral (within certain limits) when tokens are exchanged on the secondary market. Of course, we need to be sure that the risk model is adequate to the task of accounting for risk at the point of exchange.

We should also recognize that the risk model, however good, is no substitute for a vigorous counterintelligence (CI) program – in the future as now. Where the Aldrich Ames of the past insinuated himself into sensitive programs and greater access by “human engineering,” the Aldrich Ames of the future will be wheedling tokens from management, co-workers, and any secondary markets that he is authorized to access. What is different in Phase 1 is that these transactions are all trackable, auditable, and centrally reported. Far from being a substitute for CI awareness, the tokenization of risk is enabling of more and better CI.

Finally, to be clear on this point, one should not imagine that the individual government employee (or soldier in a trench) is going to be doing anything so esoteric as trading tokens on a secondary financial market, either as part of their job or in their spare time! If tokens are traded, it will be at high management levels, say between agencies, or among large project offices within an agency as a part of other resource allocation decisions. The individual employee just does his or her job: the tokens required for them to

do their work are supplied by their management chain. Most consumers of tokens will be operating with the equivalent of a "debit card," and will not be roaming freely in the marketplace. Like office furniture or computer services, risk-denominated tokens are resources that should be supplied to people as invisibly as possible, so that those people remain focused on getting their jobs done.

4.4 Originator Control is Eliminated in Phase 2

The critical changes between Phase 1 and Phase 2 are that all the different token types are collapsed to a few broad currencies (e.g. "green" tokens for tactical secrets and "gold" tokens for strategic ones), and that the function of issuing tokens is taken away from the information producers and explicitly vested in a central authority with (what amounts to) central bank authority. There are some less obvious details, however, which we now consider.

4.4.1 Tokens Collapsed to a Few Broad Token Types

The reason for having more than one token type, instead of just one, is that different kinds of secrets require different protection profiles over time. Tactical secrets require a high degree of protection "right now," but don't require long term protection. Tokens for tactical secrets ("green tokens") may therefore be issued relatively freely to localized user groups whose interests coincide with maintaining security, for example, to the soldiers in a company about to go on a mission in harm's way. Unused green tokens can likewise be collected after the mission is completed.⁷

⁷Can Corporal Z, who is actually a spy, use his green tokens to get secret tactical information about an operation on the other side of the world, in which he has no involvement? No. Although he has the right kind of tokens for tactical secrets, the risk model assigns a prohibitively large cost to any Corporal acquiring information not directly related to his unit's assignment (in this case, his local geographical area). On the other hand, can

Strategic secrets are different from tactical secrets in that they may require protection for generations, nuclear weapons design information and cryptographic methods being good examples. Strategic ("gold") tokens are the broad currency for this kind of secret. Gold tokens are distributed to organizations and individuals more selectively than green ones, and likely with fewer sharp peaks and valleys in the rate of distribution over the years. The harm represented by green token risk, while real, is over quickly. The harm represented by gold token risk may last for decades, and may not even be immediately evident.

4.4.2 Within A Broad Token Type, Access Is Now Fungible

If the number of different-colored tokens in Phase 2 is kept small (e.g., green, gold, and no others), then stovepipes controlled by the information producers disappear with the collapse of token types. IMINT, SIGINT, and HUMINT products can all be obtained in exchange for the same kind of token: green for products requiring only tactical protection; gold for products needing strategic protection.

In Phase 2, information consumers are able to optimize the use of their tokens, exchanging them for the information that they can best use. The economy is now market or "pull," not command or "push." Information consumers "vote with their tokens" for the most valuable products, and information producers can be incentivized appropriately.

We who live under the present security system may find fungibility a bit strange; it may even seem insecure. Consider the hypothetical example of a SIGINT analyst listening to a sensitive intercept in which terrorists discuss the design of an improvised nuclear weapon. Under the present system there are few SIGINT analysts with clearances that allow them access to Restricted Data (RD, that is, nuclear design data), and fewer nuclear designers

Colonel W get such information? Yes. A Colonel would be expected to need a broad range of information, and his risk factor would be constant across theaters of operation.

with compartmented SIGINT clearances. Arranging the right conversation between a SIGINT analyst and a nuclear designer may take days, weeks, or months, just for each person to be cleared by the other's security system. Under the new system, however, the analyst just calls the expert he needs and "pays" the token-denominated cost for the two of them to have a conversation spanning both SIGINT and RD. That cost might be high or low, depending on both persons' individual risk factors. The key point, however, is that the aggregate "harm" of the transaction is being correctly accounted for.

4.4.3 Phase 2 Requires Both a Risk Model and a Damage Model

A subtle point: In Phase 1, tokens were denominated by "risk"; in Phase 2, they must be denominated by "harm." (That is why we used the word "harm" rather than "risk" in the last sentence of the preceding section.)

To see why, consider two documents, one Secret (A) and the other Top Secret (B). In Phase 1, A and B have different token types (say, T010000 for A and T020000 for B). One particular type of access transaction (say, one-day, soft-copy access by a TS cleared individual) *costs the same* for A and B, because the risk is the same. This is true even though the harm from losing B is much greater than that from losing A. The way aggregate risk is limited in Phase 1 is that there are many fewer tokens of type T020000 issued than of type T010000, reflecting the lower tolerable risk for compromising the Top Secret document. So Phase 1 controls "harm" by limiting the availability of tokens, with a granularity extending in principle down to the single document level.

This no longer works in Phase 2, because both documents A and B are now available to users for the same fungible gold tokens. We must therefore charge a transaction price that reflects both the risk of the transaction and the damage, or adverse consequence, of loss to national security. The

combination of these two factors we call "harm".⁸

To implement this idea, a new factor – a "damage" factor – must be assigned to each document (or session, service, etc.). If the damage factor for the average Secret document is 1, the damage factor might be 50 for the average Top Secret document, and 10000 for a document that exposes a particularly fragile long-term intelligence capability. Just as we needed a risk model to get to Phase 1, we need a damage model to get to Phase 2; otherwise we cannot achieve the fungibility of, and user optimization among, different levels of secrets.

The cost of a transaction (access of a secret by a consumer) is given by the formula

$$\left(\begin{array}{c} \text{Cost of Transaction} \\ \text{["harm"]} \end{array} \right) = \left(\begin{array}{c} \text{Damage} \\ \text{Factor} \end{array} \right) \times \left(\begin{array}{c} \text{Risk} \\ \text{Factor} \end{array} \right)$$

4.4.4 Tokens Are Created And Distributed by a National Authority (Central Bank)

Since tokens in Phase 2 are fungible, it no longer makes sense for them to be created and distributed by the individual information-producing stovepipes. Indeed, Phase 2's *raison d'être* is to eliminate originator control as a normal way of doing business.

We therefore need a national authority whose function is to create tokens and distribute them at the highest level (i.e., in large blocks to Departments and Agencies). This national authority – which we might vest in SECDEF, or perhaps jointly between DoD, DHS, and a cabinet-level intelligence director (should one exist) – functions like a central bank. It monitors the health of the "currency" by weighing user pressure (more tokens, inflationary) against

⁸For the statistically minded reader, "damage" is the cost to the nation of a lost secret, "risk" is the (unnormalized) probability that a specified transaction will cause that loss, and "harm" is the expectation value of the cost to the nation. It is "harm" that we want to control.

security pressure (fewer tokens, deflationary). It also serves as the central collection and analysis point for audit records, and for security infraction and violation statistics. And it is the owner of the risk model.

In just the way that the Federal Reserve employs economists to study the economy, the national token-issuing authority employs security experts and user-community liaisons to study the nation's information security "economy." Some sample issues might be:

- If Department Q needs more tokens to do its job, should they be created *ab initio*, or should they be reallocated from Department R?
- In the light of a rash of espionage cases, should the supply of tokens be tightened, or do we only need to adjust some key factors in the risk model?
- How much aggregate risk would be mitigated by the adoption of some particular technology standard, and how will this be reflected in the risk model?
- Is Agency F over-valuing the damage factor for its products and thus making them too expensive, a dislocation in the market?

4.4.5 In a Tokenized System, Personnel Reliability Is a Continuous Variable

Several major espionage cases have shown a systemic weakness in the present security system, namely the fact that individuals are most often treated as either "fully trusted" (cleared) or "full untrusted" (uncleared). That is, trust is treated as a discrete, not a continuous, variable. A major reason for this is that a down-transition between these two states – revoking someone's clearance – is so drastic an action that line managers, and even security managers, try to avoid it at almost any cost.

The Aldrich H. Ames case is a particularly famous, and perhaps egregious, example of this phenomenon.⁹ Not wanting to rock the boat, managers at multiple levels dismissed, or explained away, warning signs that should have accumulated to quite a damning profile. In effect, each "explanation" reset Ames' risk odometer back to zero.

The risk model that we have posited for a tokenized system (both Phase 1 and Phase 2) provides, in effect, a continuously variable level of trust. The individual manager, therefore, is never faced with an all-or-nothing decision about whether to seek suspension of an employee's security access. Instead, the manager has clearly defined, and separable, responsibilities in functional ("getting the job done") and security ("work securely") roles.

The manager's security responsibility is simply to ensure that relevant facts known to him about an employee are also known to the risk model. The manager is not the only source of these facts, and he or she is not a "cop." But it is reasonable to require that managers report, according to objective criteria, security-related factual information about their employees. They will be more likely than now to do this knowing that the effect of such information will be incremental only, and not likely to trigger anything as drastic as the revocation of an employee's clearance (as it might today).

The manager's functional responsibility is to manage his or her unit's tokens wisely. If the risk model assigns a high risk to one particular employee, then the cost of transactions (access to classified information) for that employee will become high. The manager may make the purely economic decision to change that employee's job assignment, so that it requires less access to secret information and is less costly to the token budget. If an employee's transaction cost becomes prohibitively high, then the manager may counsel the employee either (i) to take actions that would lower that cost, or (ii) to seek other employment.

⁹The 1994 CIA IG report, available at <http://www.loyola.edu/dept/politics/intel/hitzrept.html> makes interesting reading even today.

How can an employee lower their own risk, as reflected in their token cost? Should this be allowed at all? Clearly an employee's risk should only be lowered on the basis of evidence that can not readily be faked by a spy or other malefactor.

Example: An employee purchases an expensive house for cash (as in the Ames case). This fact is assimilated automatically from public records by the risk model.¹⁰ The employee also reports the fact himself, because he knows that failure to make a timely report of such a significant event would cause his risk factor to go through the roof!

Even when reported, buying a house for cash, in and of itself, increases the employee's risk factor in the risk model. This increase is immediately propagated into all of his token transactions. He must go to his manager and ask for an increase in his token budget. The manager asks him to look into the cause of the increase.

A secure website provides the employee with personalized information about possible actions that he might take to reduce his risk factor. Here, the principal suggestion from the web site is that the employee should provide for examination detailed records that trace the origin of the funds used in the house purchase.

The employee furnishes copies of his late, rich aunt's death certificate, and of her will. He is also able to furnish his aunt's brokerage statements showing that her assets go back 20 years or more. *These records are field-checked* by DoD/DSS investigators – the ones who are now wasting their efforts on routine security investigations, rather than on (as here) targeted verification of important facts.

The employee's risk factor returns *almost* to what it was before. Almost, but not quite: the risk model recognizes a small

¹⁰It is a condition of employment that the employee has agreed that such public records can be collected.

residual risk that the evidence might have been faked, despite the documentation submitted. Risk must be accounted for!

And what if the documents didn't check out, and were judged to have been forged? In the purest model, the individual continues to be employed, but his risk factor is now a billion times (say) what it was previously. He can't afford to see a single classified document! In practice, of course, certain thresholds trigger qualitatively different actions – such as arrests for fraud.

5 SUMMARY AND CONCLUSIONS

Although we have given considerable attention to a very specific, and rather futuristic, proposal (Section 4), we do not want its exotic detail to eclipse our more general findings and conclusions, which we enumerate here.

1. The present system of classification, clearances, and access protection is broken. It is giving rise to a complicated, and largely uncontrolled, set of workarounds. The perceptual gap between those who use information and those who protect it is insupportably wide.

2. There could be immediate benefits from greater risk acceptance by technical solutions, in exchange for a decrease in even riskier human behaviors. We have gone too far in declining to certify technology solutions that have small weaknesses, even when those weaknesses are exploitable only by the most elite of our peer adversaries and only rarely at that. We should not demand perfection of technical solutions.

3. "Enclave" or "Community of Interest" concepts should be developed into actual pilot programs that would allow classified information to be utilized, under carefully bounded circumstances, by users who are so authorized but do not necessarily hold clearances for the information.

4. The long term path forward should de-emphasize the concept of a "trusted" person, and instead emphasize the concept of transactional risk associated with *any* individual's access to *any* secret. That risk depends on attributes of the individual (including his or her clearance level), but it also depends on the nature of the transaction. Technologies should be developed to offer lower-risk transaction types (e.g., tamper-proof, soft-copy display) that can be traded off against higher-risk users (e.g., uncleared soldiers in the field).

5. As an additional long-term goal, risk (or, more precisely, "harm") should be managed in aggregate at the national level, and not subject to the

vagaries of uneven management by individual information producers. We have suggested an economic model for managing "harm," based on freely-exchangeable "tokens." That model may seem too extreme to some. Maybe it is! However, the underlying three principles on which it is based ought to be touchstones of any system that aims at achieving the full value of secret information:

- *Measure* risk.
- *Set* an acceptable risk level.
- *Require* distribution all the way up to the acceptable risk level.